Course Title: Artificial Intelligence **Sheet #4**

The file ex2data1.txt contains the dataset for our sheet.

- In the previous sheet, you found the optimal parameters of a linear regression model by implementing gradient descent. Repeat the same assignment, however this time, instead of taking gradient descent steps, you will use an Matlab built-in function called fminunc, to find the best parameters θ for the logistic regression cost function.

$$J(\theta) = \frac{1}{m}\sum_{i=1}^{m}\left[-y^{(i)}\log(h_\theta(x^{(i)})) - (1-y^{(i)})\log(1-h_\theta(x^{(i)}))\right],$$

and the gradient of the cost is a vector of the same length as θ where the j^{th} element (for $j = 0, 1, \ldots, n$) is defined as follows:

$$\frac{\partial J(\theta)}{\partial \theta_j} = \frac{1}{m}\sum_{i=1}^{m}(h_\theta(x^{(i)}) - y^{(i)})x_j^{(i)}$$

- After learning the parameters using gradient decent or fminunc, use the model to predict whether a particular student will be admitted. For a student with an Exam 1 score of 45 and an Exam 2 score of 85, you should expect to see an admission probability of 0.776. your task is to write a function predict.m. The predict function will produce "1" or "0" predictions given a dataset and a learned parameter vector θ.

*p = PREDICT(theta, X) computes the predictions for X using a threshold at 0.5 (i.e., if sigmoid(theta'*x) >= 0.5, predict 1)*

- Implement regularized logistic regression to predict whether microchips from a fabrication plant passes quality assurance (QA). During QA, each microchip goes through various tests to ensure it is functioning correctly. Suppose you are the product manager of the factory and you have the test results for some microchips on two different tests. From these two tests, you would like to determine whether the microchips should be accepted or rejected. To help you make the decision, you have a dataset of test results on past microchips, from which you can build a logistic regression model. **The file ex2data2.txt contains the dataset for our sheet.**
- First load the data and display it on a 2-dimensional plot.
- Implement the cost function and gradient for regularized logistic regression, where the cost function is

$$J(\theta) = \left[-\frac{1}{m}\sum_{i=1}^{m} y^{(i)}\log\left(h_\theta(x^{(i)}) + (1-y^{(i)})\log 1 - h_\theta(x^{(i)})\right)\right] + \frac{\lambda}{2m}\sum_{j=1}^{n}\theta_j^2$$

- Use the final θ values obtained by the gradient decent algorithm to plot the decision boundary on the training data.
- Try out different regularization parameters for the dataset to understand how regularization prevents overftting (Try λ=1, 0 and 100) and specify the fitting case with each trail.

Best wishes

Dr. Sherin El Gokhy